

GPTKB 2.0: Browsing, Querying, and Auditing a Disambiguated LLM-Derived Knowledge Base

Yujia Hu¹ Tuan-Phong Nguyen² Simon Razniewski¹

¹ScaDS.AI Dresden/Leipzig & TU Dresden, Germany

²VNU University of Engineering and Technology, Hanoi, Vietnam

{yujia.hu, simon.rzniewski}@tu-dresden.de tuanphong@vnu.edu.vn

Abstract

We present a web demo for exploring a large-scale disambiguated knowledge base (KB) materialized from a large language model (LLM). GPTKB 2.0 contains 38.4M triples over 1.6M canonical entities, together with 207.6K consolidated relations and 66K consolidated classes. Unlike prior LLM-derived knowledge bases that largely identify entities by surface strings, GPTKB 2.0 performs context-guided disambiguation during recursive KB construction, separating homonyms and merging synonymous mentions as facts are elicited. The demo makes this process inspectable: users can browse entities, follow links across the KB, and audit the provenance of individual facts, including surface forms, candidate matches, source triples, and disambiguation decisions. The interface further supports structured SPARQL queries, natural-language questions translated to SPARQL, and entity linking from user-provided text to canonical GPTKB 2.0 entries. GPTKB 2.0 is available at <https://gptkb.org/>, with the full KB downloadable for offline use.

1 Introduction

General-domain KBs (Vrandečić and Krötzsch, 2014; Auer et al., 2007; Suchanek et al., 2007) are important and longstanding backbones for AI applications, and entity disambiguation recognized as a core challenge in KB construction. More recently, large language models (LLMs) have emerged as implicit repositories of factual knowledge (Petroni et al., 2019), inspiring a line of work that constructs knowledge bases directly from LLMs by materializing their parametric knowledge in structured form (Nguyen et al., 2024; Cohen et al., 2023; Hu et al., 2025, 2026a; Parović et al., 2025). However, these approaches largely rely on surface strings as entity identifiers. Yet surface strings are inadequate identifiers, exhibiting two complementary failure modes: without further triple context, they




Figure 1: Budapest in GPTKB 2.0.

conflate *homonymous* entities that share identical entity label (e.g., *Munich* in [*Europe*, *hasMajorCity*, *Munich*] vs. *Munich* in [*Mathieu Amalric*, *notable-Work*, *Munich*]) and fragment *synonymous* strings that denote the same entity (e.g., *New York City* and *The Big Apple*).

We present GPTKB 2.0, a disambiguated, general-domain KB built entirely from an LLM, together with a web interface that makes both the knowledge and its provenance explorable. GPTKB 2.0 comprises 38.4M triples over 1.6M canonical-

ized entities, 207.6K consolidated relations, and 66K consolidated classes. Disambiguation is performed on the fly during construction, driven by a simple observation: the eliciting triple, together with textual descriptions of existing entities, typically provides sufficient context to distinguish new entities from known ones without external resources.

The demo’s distinguishing feature is transparency. Through the interface, users can (i) browse and query the full KB via SPARQL or natural language; (ii) inspect the derivation and consolidation provenance behind any fact: all its trajectories from the seed entity, and the surface forms, candidate matches, and context considered during consolidation, so that both homonymy and synonymy resolution become auditable; and (iii) link named entities in their own input text to canonical GPTKB 2.0 entries. To our knowledge, GPTKB 2.0 is the first demo to combine disambiguation provenance, SPARQL, and natural-language querying over an LLM-derived KB.

2 Construction

GPTKB 2.0 methodology (Hu et al., 2026b) represents entities, relations, and classes as first-class elements with unique identifiers, so that disambiguation can operate across homonymous and synonymous mentions. Following curated KBs such as Wikidata (Vrandečić and Krötzsch, 2014), besides unique ID, each element also carries a textual description, which serves as the context signal for disambiguation. An example for  **Budapest** is shown in Figure 1. Starting from a seed entity, GPTKB 2.0 is built by a recursive, context-guided pipeline that expands the KB along the subject–object frontier, disambiguating each new mention on the fly. The pipeline proceeds through four context-guided steps: elicitation, named entity recognition (NER), and disambiguation.

Elicitation and NER Both steps rest on one principle: a bare label is ambiguous, so we condition the LLM on context. Given an entity, the LLM is prompted for triples; since facts about *Munich* could concern the city or the 2005 film, we include the entity’s description so that elicited facts pertain to the intended entity. Each elicited object is then classified as a literal or a named entity, *01099* is a literal in *[Neustadt Dresden, postal-Code, 01099]* but an entity in *[Frisch, partOfBand, 01099]*—using the source triple as context.

Entities	1,592,185
Triples	38,450,135
Relations	207,633
Classes	66,523
Triple objects (entities)	15.8M
Triple objects (literals)	22.6M
Avg. triples per entity	38.4
Avg. outlinks per entity	15.8
Entities novel to Wikidata*	36.8%

* Validated on 1,000 samples.

Table 1: Overall statistics of GPTKB 2.0.

Disambiguation Named-entity objects are resolved against the KB to merge synonyms and separate homonyms. The LLM is prompted with the target entity label, its source triple, and the labels and descriptions of the nearest candidates by embedding similarity; it either matches an existing entity, marks it new, or abstains. A match inherits the incoming label as an alias. For instance,  **Munich** (city) and  **Munich** (film) are kept distinct through their differing source-triple context; *The Big Apple* is mapped to  **New York** based on the source triple *[Emerald City, contrastsWith, The Big Apple]*.

3 Overview

As shown in Table 1, GPTKB 2.0 contains 38.4M triples and 1.6M disambiguated entities, together with 207.6K relations and 66.5K classes.

4 Web Interface

GPTKB 2.0 is served at <https://gptkb.org/>, which provides a user interface for keyword-based entity search and link-based exploration to discover connected entities. The interface is built with the Django framework and served via Nginx, with the KB stored in an OpenLink Virtuoso triple store. To our knowledge, this is the first demo interface to make the consolidation process fully traceable while combining SPARQL querying, natural-language (text-to-SPARQL) querying, and entity linking over an LLM-derived KB.

4.1 Access

Entities can be reached through several routes:

1. The start page links directly to featured entities such as *Vannevar Bush* and *San Francisco*.
2. A search field in the top-right corner supports string-based lookup.

Hyde Park E71393

public park royal park urban park

Hyde Park is one of London's largest and most famous royal parks, known for its expansive green spaces, recreational activities, and historic landmarks like Speakers' Corner and the Serpentine.

Copy ID Copy link Try in SPARQL Jump to:

Surface forms Disambiguation Statements Elicitation

Referenced by

NED1 Entity disambiguation (via context triple) gpt-5-mini-2025-08-07

Target entity: Hyde Park
Context triple: [Mayfair, borderedBy, Hyde Park]

A. **Hyde Park**
Hyde Park is a historic central public park in Sydney, Australia, known for its tree-lined avenues, monuments, and recreational green space.

B. **Hyde Park**
Hyde Park is a historic town in New York's Hudson Valley best known as the longtime home of President Franklin D. Roosevelt and several notable Gilded Age estates.

C. **Hyde Park**
"Hyde Park" is a Caroline-era comedy play by James Shirley, known for its witty depiction of London society and romantic intrigues set around the famous London park.

D. **Hyde Park**
Hyde Park is a historic and upscale residential and shopping district located just southwest of downtown Tampa, Florida.

E. **Hyde Park**
Hyde Park is a historic neighborhood on Chicago's South Side known for its cultural diversity, lakefront parks, and institutions like the University of Chicago and the Museum of Science and Industry.

F. **None of above.** chosen

G. **Unsure** - the case is ambiguous/there is not enough information to decide.

Figure 2: *Hyde Park* in London is recognized as distinct from other synonymous entities.

- When an entity's ID is known, its page can be accessed directly at <https://gptkb.org/entity/<ID>> (e.g., E14 for the 🇺🇸USA).

4.2 Traceable decision-making

Each entity's page exposes the disambiguation decisions behind it, making both homonymy and synonymy auditable. To trace homonymy handling, users can inspect the candidates presented to the LLM when this entity was disambiguated, revealing how it was kept distinct from same-label candidates and judged new to the KB at the point it first appeared. As shown in Figure 2, clicking **Show** under *How this entity was disambiguated* reveals the process by which 🇺🇸Hyde Park in London is separated from its homonyms existing in GPTKB 2.0. To trace synonymy handling, the surface forms consolidated into this entity are listed as its aliases; clicking one opens the prompt for the corresponding disambiguation step, showing how that form was merged. As shown in Figure 3, the alias list of 🇺🇸New York City lets users click an alias (here, *The Big Apple*) to reveal the mentions mapped under that label; the info button then opens the prompt used at that step, listing the target entity, its source triple, and the retrieved candidates.

4.3 Entity Linking

Beyond showing how entities within GPTKB 2.0 were linked during construction, the interface integrates LELA (Haffoudhi et al., 2026a,b), an entity-linking system, over GPTKB 2.0. A user enters text and clicks **Link entities**; LELA detects named-entity mentions and links each to its corresponding entity in GPTKB 2.0, or marks it NIL when no match exists.

GPTKB v2 Explore About

Entity linking

Paste any text — each mention is linked to its GPTKB v2 entity, disambiguated by context (the same word can resolve to different entities).

Mercury is the closest planet to the Sun, yet in Roman mythology Mercury was the winged messenger of the gods.

Link entities Try: Mercury x2 Cambridge x2 Java x2 Paris x2 Apple x2 Jordan x2 Gandhi 110 / 2000

Linked text 2 mentions · 2 linked

Mercury is the closest planet to the Sun, yet in Roman mythology Mercury was the winged messenger of the gods.

linked opens the entity page (hover for a preview) · NIL no DisambKB match (click to see the rejected candidates)

Each decision is auditable through a disambiguation-evidence panel that lists every candidate the linker weighed. Candidates are retrieved in two ways: surface form matches, ranked by how often each entity has been observed under that surface form, and a BM25 text-relevance fallback that fills the remaining slots. Notably, context can override raw frequency: although the planet is the most frequent entity linked to *Mercury* (observed 53 times), the mythological mention is correctly resolved to the Roman deity (observed only 16 times), which the surrounding context favors. This transparency lets

New York City E40

city municipality

Top #9 by total degree (32,631 = 32,557 in + 74 out)

New York City is the largest city in the United States, a global center of finance, culture, media, and technology.

Copy ID Copy link Try in SPARQL Jump to: Surface forms

Statements Elicitation Referenced by

All labels observed (58)

Label	Occurrences
New York City CANONICAL	28,933
New York City, New York, United States	1,758
The Big Apple	4

Disambiguation evidence for New York City via surface form

"The Big Apple"

Triples (4)

Triples where some other subject referred to this entity as "The Big Apple".

New York City → nickname → "The Big Apple"
↳ resolves to New York City

Emerald City → contrastsWith → "The Big Apple"
↳ resolves to New York City

Empire State → hasRelatedNickname → "The Big Apple"
↳ resolves to New York City

The Big Guava → echoesNickname → "The Big Apple"
↳ resolves to New York City

NED1 Entity disambiguation (via context triple)

gpt-5-mini-2025-08-07

Target entity: The Big Apple
Context triple: [Emerald City, contrastsWith, The Big Apple]

A. Manhattan
Manhattan is the densely populated, iconic core borough of New York City, known for its skyscrapers, cultural institutions, and role as a global financial and media center.

B. New York City **chosen**
New York City is the largest city in the United States, a global center of finance, culture, media, and technology.

C. City of Big Shoulders
City of Big Shoulders is a poetic nickname for Chicago that evokes its history as a powerful industrial and working-class metropolis.

D. Empire State
Empire State is a well-known nickname for the U.S. state of New York, reflecting its historical wealth, influence, and prominence.

E. Brooklyn
Brooklyn is a populous and culturally diverse borough of New York City known for its distinct neighborhoods, arts scene, and iconic landmarks like the Brooklyn Bridge.

F. None of above.

G. Unsure - the case is ambiguous/there is not enough information to decide.

Figure 3: The Big Apple is mapped to existing entity *New York City*.

users see not only which entity was chosen but which alternatives were considered and rejected.

1 Mercury person → Mercury E18768

Mercury is the smallest and innermost planet in our Solar System, known for its extreme temperature variations and heavily cratered surface.

SURFACE-FORM MATCHES

ranked by observed frequency (×N = seen N× as this surface form)

1 Mercury E18768	✓ CHOSEN	×53
2 Mercury E6367		×47
3 Mercury E397485		×16
4 Mercury E1253343		×1
5 Mercury E882902		×1

BM25 FALLBACK text-relevance score — fills slots no surface form matched

6 Mercury Topaz LS E1082172	15.2
7 Mercury Mystique E741113	14.4
8 Mercury Capri E343146	14.2
9 Ford Mercury models (selected years) E250929	14.1
10 Mercury Monarch E754360	14.1

Show all 15

2 Mercury person → Mercury E397485

Mercury is the Roman god of commerce, communication, and travel, often depicted as a swift messenger of the gods.

SURFACE-FORM MATCHES

ranked by observed frequency (×N = seen N× as this surface form)

1 Mercury E18768		×53
2 Mercury E6367		×47
3 Mercury E397485	✓ CHOSEN	×16
4 Mercury E1253343		×1
5 Mercury E882902		×1

BM25 FALLBACK text-relevance score — fills slots no surface form matched

6 Mercury Topaz LS E1082172	15.2
7 Mercury Mystique E741113	14.4
8 Mercury Capri E343146	14.2
9 Ford Mercury models (selected years) E250929	14.1
10 Mercury Monarch E754360	14.1

Show all 15

4.4 SPARQL Query

Because GPTKB 2.0 is a fully materialized, structured KB, its contents can be queried directly rather than through repeated prompting of the underlying model. The KB is stored in an OpenLink Virtuoso triple store and exposed through a SPARQL endpoint at <https://gptkb.org/query/>, supporting structured queries, filtering, and aggregation over all 38.4M triples. This turns questions about the model’s knowledge that would otherwise require bespoke, large-scale prompting (e.g., thousands of templated probes (Kotek et al., 2023)) into single declarative queries evaluated with mature database technology. In the following, we show enabled analyses.

Most frequent classes The most frequent classes in GPTKB 2.0 can be queried via:

PREFIX gptkbp: <<https://gptkb.org/prop/>>

```
SELECT ?o (COUNT(*) AS ?freq)
WHERE {
  ?s gptkbp:P0 ?o .
}
GROUP BY ?o
ORDER BY DESC(?freq)
LIMIT 20
```

label	freq
human	166328
municipality	66232
tourist attraction	51891
fictional beaver	29751
song	25408
neighborhood	24317
building	23436
film	23376

Entities with most aliases Entities with most aliases in GPTKB 2.0 can be queried via:

```
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX skos: <http://www.w3.org/2004/02/skos/core#>

SELECT ?e ?eLabel (COUNT(?form) AS ?nforms)
WHERE {
  ?e (rdfs:label|skos:altLabel) ?form
}
GROUP BY ?e ?eLabel
ORDER BY DESC(?nforms)
LIMIT 20
```

e	eLabel	nforms
gptkb:E14	United States of America	128
gptkb:E11841	Academy Award for Best Cinematography	109
gptkb:E1331	UNESCO World Heritage Site	93
gptkb:E732	United Kingdom	88
gptkb:E2390	Cannes Film Festival	81
gptkb:E1004	National Historic Landmark	80
gptkb:E45485	Minister of Defence	77

As shown above, 128 distinct surface labels map to the entity *United States of America*. This illustrates the value of resolving entities during construction rather than afterward: because each merge is available to subsequent elicitation steps, later facts are gathered against the consolidated entity, avoiding the fragmentation that a post-hoc pass would first have to detect and then repair.

4.5 Query in Natural Language

For users unfamiliar with SPARQL, the demo integrates GRASP (Walter and Bast, 2025), an agent that answers plain-English questions over GPTKB 2.0. Given a question, the agent iteratively explores GPTKB 2.0 through tool calls, resolving the entities, properties, and classes it needs, then composes and executes a SPARQL query to produce the answer. For example, “*What is Budapest known for?*” is resolved by locating the *Budapest* entity (E13406) and the *notableFor* property (P22), yielding a concise list of answers.

 GPTKB v2 [Explore](#) [About](#)

Ask GPTKB v2

Ask a question in plain English — an LLM agent explores GPTKB v2, writes a SPARQL query that answers it, and shows you both the result and the query.

What is Budapest known for?

 Ask

Try:

[Notable for](#)

[Count](#)

[Surface form](#)

[Join](#)

Result

Budapest is known for its Danube River panorama, historic architecture, ruin pubs, spa culture, and thermal baths.

Got 5 rows and 1 column in 0.00s:

thing
"Danube River panorama"
"historic architecture"
"ruin pubs"
"spa culture"
"thermal baths"

SPARQL query [Open in query editor](#) [Copy](#)

```
PREFIX gptkb: <https://gptkb.org/prop/>
PREFIX gptkb: <https://gptkb.org/entity/>
SELECT ?thing WHERE { gptkb:E13406 gptkb:P22 ?thing }
```

KB items the agent used [Hide items](#)

The GPTKB v2 entities, properties and concepts the agent resolved for the final query — each with its KB description, linking to its page.

ENTITIES

Budapest E13406

Budapest is the capital and largest city of Hungary, renowned for its historic architecture, thermal baths, and prominent location along the Danube River.

PROPERTIES

notableFor P22

Indicates that an entity is especially recognized or distinguished for a particular quality, achievement, characteristic, or role.

Consistent with the rest of the interface, every step is transparent: the demo exposes the agent’s full reasoning trace, the generated SPARQL query (which users can open in the query editor), and the KB entries the agent resolved, each linked to its entity page.

Agent trace [Hide trace](#)

9 steps · 15s

The agent is GRASP (Walter & Bast, ISWC 2025), self-hosted against GPTKB v2’s RDF export. Your question goes to it (openai/gpt-oss-120b) — its only knowledge of GPTKB v2 comes through tool calls, and every step it takes streams here, live. Its loop:

Question sent

Thinking

We need to answer: “What is Budapest known for?” Likely they want notable things about Budapest. In knowledge graph, may have statements like <entity> gptkb:P... perhaps “known for” property. We need to find property representing “known for” maybe label includes “known for”. Let’s search properties index with query “known for”.

Search “known for”

Results (page 1 of 7): 1. namedForKnownFor (gptkb:P95239), also known as “P95239” 2. fameFor (gptkb:P52554), also known as “P52554” 3. bekanntFür (gptkb:P161540), also known as “P161540” 4. nowKnownFor (gptkb:P143719...

[Show more](#)

Thinking

We need the entity for Budapest. Find entity.

Search “Budapest”

Results (page 1 of 9): 1. Budapest (gptkb:E13406), also known as “Budapest city centre”, “Budapest Municipality”, “Budapest with Tatabánya”, “Budapest metropolitan area”, “Municipality of Budapest”, etc. 2. Budapest Metr... [Show more](#)

Thinking

Now we need property linking Budapest to known for. Likely property “fameFor” or “notableFor”. Let’s check usage via list where subject = Budapest and property = fameFor etc.

List “gptkb:E13406”

No triples (page 1 of 1)

Thinking

Try notableFor (P22).

List “gptkb:E13406”

Triples (page 1 of 1): | subject | property | object | | Budapest (gptkb:E13406) | notableFor (gptkb:P22) | “Danube River panorama” | | Bud... [Show more](#)

Thinking

Great. So answer: Budapest is known for Danube River panorama, historic architecture, ruin pubs, spa culture, thermal baths. Could phrase “Budapest is known for its thermal baths, historic architecture, ruin pubs, spa culture, and Danube River scenery.” Need to produce answer via answer function with final SPARQL query. Provide query that returns these objects. So final query: SELECT ?thing WHERE ... [Show more](#)

5 Evaluation

5.1 Disambiguation

Reliable link-based exploration depends on correct disambiguation: linked entities must be neither wrongly merged nor wrongly split. We manually assess both error directions on samples of 100 cases each (Table 2). **Merge precision** measures whether entities combined into one are truly the same: 98%

Judge	Human	LLM	LLM ⁻
NED precision			
Merge correct	94.5%	—	—
Split correct	97.5%	—	—
Triple precision			
True	94.5%	92.8%	80.0%
Plausible	2.5%	5.5%	10.3%
Implausible	1.0%	1.2%	2.0%
False %	2.0%	0.5%	7.7%
Entity factuality			
Verifiable	96.0%	92.3%	—
Plausible	2.0%	6.7%	—
Unverifiable	2.0%	1.0%	—

Table 2: Manual and automatic evaluation of GPTKB 2.0. Human evaluation uses $n = 400$ for NED precision and $n = 200$ for triple precision and entity factuality; LLM evaluation uses $n = 1,000$.

of same-label objects merged into an existing entity are correct (homonymy), as are 91% of differently-labeled aliases resolved to an existing entity (synonymy); false merges (9%) in the latter are the main error mode. **Split precision** measures whether entities kept separate are truly distinct: all 100 same-label pairs assigned distinct IDs are genuinely different (homonymy), and for 95% of entities no synonymous duplicate appears among their top-20 label-embedding neighbors (synonymy; an approximation, since duplicates with dissimilar surface forms may be missed). The pipeline thus handles both homonymy and synonymy reliably in both directions, with false synonym merges the largest residual error.

5.2 Overall Quality of GPTKB 2.0

Since exhaustive evaluation of a 38M-triple KB is infeasible, we assess randomly sampled subsets. Following YAGO (Suchanek et al., 2007) and recent factuality metrics such as FactScore (Min et al., 2023) and VeriScore (Song et al., 2024), we verify samples against web-retrieved evidence, using two paradigms: 200 items validated by humans and 1,000 items judged by an agentic pipeline. Precision exceeds 90% for both triples and entities (Table 2). For **triples** (labeled true, plausible, implausible, or false), 94.5% are true and 2.0% false on the human sample, and 92.8% true and 0.5% false on the automatic sample. In an ablation where the judge sees each triple without the subject entity’s description, the true rate drops to 80%, underscoring how much descriptions contribute to interpretability. For **entities** (labeled verifiable, plausible, or unverifiable), 96% are verifiable on the

human sample and 92.3% on the automatic sample.

6 Related Work

Disambiguated Knowledge Bases Managing entity identity is a longstanding challenge in KB construction. Curated KBs (DBpedia (Auer et al., 2007), YAGO (Suchanek et al., 2007), Wikidata (Vrandečić and Krötzsch, 2014)) inherit canonicalization from human-maintained resources, and mention-grounding methods reuse such inventories (Hoffart et al., 2011), but both bound coverage to what those resources contain; text-based systems (Carlson et al., 2010) broaden coverage yet often set consolidation aside. LLM-derived KBs (Cohen et al., 2023; Hu et al., 2025; Parović et al., 2025) are tied to neither a fixed corpus nor a closed schema, but inherit no external identifiers and must disambiguate from scratch.

Traceable KB Interfaces Public KBs are typically exposed through browsers, entity pages, and SPARQL endpoints that present the KB as a finished artifact showing what it contains, not how each entry arose. For LLM-derived KBs, whose entries are generated and consolidated automatically, the provenance behind each fact and disambiguation decision is itself evidence of trustworthiness. GPTKB 2.0’s interface makes this process traceable, surfacing for each fact its derivation trajectory, the candidates weighed during disambiguation, and the context behind each decision.

7 Conclusion

We presented GPTKB 2.0, a disambiguated knowledge base materialized entirely from LLM, together with a web interface that makes both its knowledge and the disambiguation decisions behind it explorable. The KB comprises 38.4M triples over 1.6M entities, 207.6K relations, and 66K classes, all consolidated on the fly during a recursive construction pipeline. The distinguishing feature of this demo is transparency: every fact carries provenance recording the surface forms, candidate matches, and context behind its disambiguation. Through the interface, users can browse and trace entities, query the KB in SPARQL or natural language, and link named entities in their own text against it. By turning an LLM’s latent parametric knowledge into a structured, queryable, and inspectable resource, GPTKB 2.0 offers both a usable general-domain KB and a step toward making the knowledge inside LLMs open to scrutiny.

References

- Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary G. Ives. 2007. [DBpedia: A nucleus for a web of open data](#). In *ISWC*.
- Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R. Hruschka Jr., and Tom M. Mitchell. 2010. [Toward an architecture for never-ending language learning](#). In *AAAI*.
- Roi Cohen, Mor Geva, Jonathan Berant, and Amir Globerson. 2023. [Crawling the internal knowledge-base of language models](#). In *Findings of EACL*.
- Samy Haffoudhi, Fabian M. Suchanek, and Nils Holtenberger. 2026a. [LELA: An end-to-end LLM-based entity linking framework with zero-shot domain adaptation](#). In *IJCAI-ECAI*.
- Samy Haffoudhi, Fabian M. Suchanek, and Nils Holtenberger. 2026b. [LELA: an LLM-based entity linking approach with zero-shot domain adaptation](#). In *ISWC*.
- Johannes Hoffart, Mohamed Amir Yosef, Ilaria Bordino, Hagen Fürstenau, Manfred Pinkal, Marc Spaniol, Bilyana Taneva, Stefan Thater, and Gerhard Weikum. 2011. [Robust disambiguation of named entities in text](#). In *EMNLP*.
- Yujia Hu, Tuan-Phong Nguyen, Shrestha Ghosh, Moritz Müller, and Simon Razniewski. 2026a. [GPTKB v1.5: A massive knowledge base for exploring factual LLM knowledge](#). In *AAAI Demos*.
- Yujia Hu, Tuan-Phong Nguyen, Shrestha Ghosh, and Simon Razniewski. 2025. [Enabling LLM knowledge analysis via extensive materialization](#). In *ACL*.
- Yujia Hu, Tuan-Phong Nguyen, and Simon Razniewski. 2026b. [GPTKB 2.0: Direct construction of disambiguated knowledge bases from large language models](#). *Preprint*, arXiv:2608.03729.
- Hadas Kotek, Rikker Dockum, and David Sun. 2023. [Gender bias and stereotypes in large language models](#). In *ACM CI*.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FActScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *EMNLP*.
- Tuan-Phong Nguyen, Simon Razniewski, and Gerhard Weikum. 2024. [Cultural commonsense knowledge for intercultural dialogues](#). In *CIKM*.
- Marinela Parović, Ze Li, and Jinhua Du. 2025. [Generating domain-specific knowledge graphs from large language models](#). In *Findings of ACL*.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) In *EMNLP-IJCNLP*.
- Yixiao Song, Yekyung Kim, and Mohit Iyyer. 2024. [VeriScore: Evaluating the factuality of verifiable claims in long-form text generation](#). In *Findings of EMNLP*.
- Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. [Yago: a core of semantic knowledge](#). In *WWW*.
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: a free collaborative knowledgebase](#). *Commun. ACM*.
- Sebastian Walter and Hannah Bast. 2025. [GRASP: generic reasoning and SPARQL generation across knowledge graphs](#). In *ISWC*.